arXiv:1502.06644v1 [stat.ML] 23 Feb 2015

On The Identifiability of Mixture Models from Grouped Samples

Robert A. Vandermeulen RVDM@UMICH.EDU
Clayton D. Scott CLAYSCOT@UMICH.EDU
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109

Abstract

Finite mixture models are statistical models which appear in many problems in statistics and machine learning. In such models it is assumed that data are drawn from random probability measures, called mixture components, which are themselves drawn from a probability measure $\mathscr{P}$ over probability measures. When estimating mixture models, it is common to make assumptions on the mixture components, such as parametric assumptions. In this paper, we make no assumption on the mixture components, and instead assume that observations from the mixture model are grouped, such that observations in the same group are known to be drawn from the same component. We show that any mixture of m probability measures can be uniquely identified provided there are $2m-1$ observations per group. Moreover we show that, for any m, there exists a mixture of m probability measures that cannot be uniquely identified when groups have $2m-2$ observations. Our results hold for any sample space with more than one element.

Keywords: Mixture Model, Latent Variable Model, Identifiability, Hilbert Space, Tensor Product

1. Introduction

A finite mixture model is a probability law based on a finite number of probability measures, $\mu_1, \ldots, \mu_m$, and a discrete distribution $w_1, \ldots, w_m$. A realization of a mixture model is first generated by first generating a component at random k, $1 \leq k \leq m$, and then drawing from μ_k. A mixture model can be associated with a probability measure on probability measures, which we denote $\mathscr{P}$. Mixture models are used to model data throughout statistics and machine learning.

A primary theoretical question concerning mixture models is identifiability. A mixture model is said to be identifiable if no other mixture model (of equal or lesser complexity) explains the distribution of the data. Some previous work on identifiability considers the situation where the observations are drawn iid from the mixture model, and conditions on $\mu_1, \ldots, \mu_m$ are imposed, such as Gaussianity (Dasgupta and Schulman, 2007; Anderson et al., 2014). In this work we make no assumptions on $\mu_1, \ldots, \mu_m$. Instead, we assume the observations are grouped, such that realizations from the same group are known to be iid from the same component. We call these groups of samples "random groups." We define a random group to be a random collection $\mathbf{X}_i$, where $\mathbf{X}_i = X_{i,1}, \ldots, X_{i,n} \overset{iid}{\sim} \mu_i$ and $\mu_i \overset{iid}{\sim} \mathscr{P}$.

Consider the set of all mixtures of probability measures which yield the same distribution over the random groups as does $\mathscr{P}$. If some element of this set other than $\mathscr{P}$ has no more components than $\mathscr{P}$ then $\mathscr{P}$ is not identifiable. In other words, there is no way to differentiate $\mathscr{P}$ from another model of equal or lesser complexity. Fortunately, with a sufficient number of samples in each

© R.A. Vandermeulen & C.D. Scott.

random group, $\mathscr{P}$ becomes the most simple model which describes the data. In this paper we show that, for any sample space, any mixture of probability measures with m components is identifiable when there are $2m-1$ samples per random group. Furthermore we show that this bound cannot be improved, regardless of sample space.

1.1. Applications of Probability Measures over Probability Measures

Though a somewhat mathematically abstract object, probability measures over spaces of probability measures arise quite naturally in many statistical problems. Any application which use mixture models, for example clustering, is utilizing a probability measure over probability measures. Moreover mixture models are a subset of a larger class of models known as latent variable models. One problem in latent variable models which has seen significant interest recently is topic modeling. Topic modelling is concerned with the extraction of some sort of topical structure from a collection of documents. Many popular methods for topic modelling assume that each document in question has a latent variable representing a "topic" or a random convex combination of topics which determines the distribution of words in that document (Blei et al., 2003; Anandkumar et al., 2014; Arora et al., 2012).

Another statistical problem which often utilizes a probability measure over probability measures is transfer learning. In transfer learning one is interested in utilizing several different but related training datasets (perhaps a collection of datasets which correspond to different patients in a study) to construct some sort of classifier or regressor for another different but related testing dataset. There are many approaches to this problem but one formulation assumes that each dataset is generated from a random probability measure and each random measure is generated from a fixed probability measure over probability measures (Blanchard et al., 2011; Maurer et al., 2013).

Finally sometimes we would like to perform statistical techniques directly on a space of probability measures. Examples of this include detection of anomalous distributions (Muandet and Schölkopf, 2013) and distribution regression (Póczos et al., 2013; Szabo et al., 2014).

1.2. How Does Group Size Affect Consistency?

Many of the applications above assume a model similar to the one we described in the first paragraph. They assume there exists some probability measure, $\mathscr{P}$, over a space of probability measures from which we have observed groups of data $\mathbf{X}_1, \ldots, \mathbf{X}_N$ with $\mathbf{X}_i = X_{i,1}, \ldots, X_{i,M_i} \overset{iid}{\sim} p_i$ and $p_i \sim \mathscr{P}$. For example in topic modeling each $\mathbf{X}_i$ is a document which contains M_i words and in transfer learning $\mathbf{X}_i$ is one of the several different training datasets. Proposed algorithms for solving these problems often contain some sort of consistency result and these results typically require that $N \to \infty$ and either $M_i \to \infty$ for all i or that $\mathscr{P}$ satisfies some properties which makes $M_i \to \infty$ unnecessary. When considering such results one may wonder what sort of statistical penalty we incur from fixing $M_i = C$ for all i.

While this question is clearly interesting from a theoretical perspective it has a couple of important practical implications. Firstly it is not uncommon for C to be restricted in practice. An example of this is topic modelling of Twitter documents, where the restricted character count keeps each M_i quite small. The second important practical consideration is that some latent variable techniques do not utilize the full sample $\mathbf{X}_i$ and instead break down $\mathbf{X}_i$ into many pairs or triplets of samples for analysis (Anandkumar et al., 2014; Arora et al., 2012). It is important to know what, if anything, is

lost from doing this. Though we do not provide a direct answer to this question, our results seem to suggest that such techniques may significantly limit what can be known about $\mathscr{P}$.

2. Related Work

The question of how many samples are necessary in each random group to uniquely identify a finite mixture of measures has come up sporadically over the past couple of decades. The application of Kruskal's theorem (Kruskal, 1977) has been used to concoct various identifiability results for random groups containing three samples. In Allman et al. (2009) it was shown that any mixture of linearly independent measures over a discrete space or linearly independent probability distributions on $\mathbb{R}^d$ are identifiable from random groups containing three samples. In Hettmansperger and Thomas (2000) it was shown that a mixture of m probability measures on $\mathbb{R}$ is identifiable from random groups of size $2m-1$ provided there exists some point in $\mathbb{R}$ where the cdf of each mixture component at that point is distinct. The result most closely resembling our own is in Rabani et al. (2013). In that paper they show that a mixture of m probability measures over a discrete domain is identifiable with $2m-1$ samples in each random group. They also show that this bound is tight and provide a consistent algorithm for estimating arbitrary mixtures of measures over a discrete domain.

Our proofs are quite different from other related identifiability results and rely on tools from functional analysis. Other results in the same vein as ours rely on algebraic or spectral theoretic tools. Our proofs basically rely on two proof techniques. The first technique is the embedding of finite collections of measures in some Hilbert space. The second technique is using the properties of symmetric tensors over $\mathbb{R}^d$ and applying them to tensor products of Hilbert spaces. Our proofs are not totally detached from the algebraic techniques but the algebraic portions are hidden away in previous results about symmetric tensors.

3. Problem Setup

We will be treating this problem in as general of a setting as possible. For any measurable space we define δ_x as the Dirac measure at x. For $\odot$ a set, σ-algebra, or measure, we denote $\odot^{\times a}$ to be the standard a-fold product associated with that object. For any natural number k we define $[k] \triangleq \mathbb{N} \bigcap [1,k]$. Let Ω be a set containing more than one element. This set is the sample space of our data. Let $\mathcal{F}$ be a σ-algebra over Ω. Assume $\mathcal{F} \neq \{\emptyset, \Omega\}$. We denote the space of probability measures over this space as $\mathcal{D}(\Omega, \mathcal{F})$, which we will shorten to $\mathcal{D}$. We will equip $\mathcal{D}$ with the σ-algebra $2^{\mathcal{D}}$ so that each Dirac measure over $\mathcal{D}$ is unique. Define $\Delta(\mathcal{D}) \triangleq \operatorname{span}(\delta_x : x \in \mathcal{D})$. This will be the ambient space where our mixtures of probability measures live. Let $\mathscr{P} = \sum_{i=1}^m \delta_{\mu_i} w_i$ be a probability measure in $\Delta(\mathcal{D})$. Let $\mu \sim \mathscr{P}$ and $X_1, \ldots, X_n \overset{iid}{\sim} \mu$. We will denote $\mathbf{X} = (X_1, \ldots, X_n)$.

We will now derive the probability law of $\mathbf{X}$. Let $A \in \Omega^{\times n}$, we have

$$\begin{aligned} \mathbb{P}(\mathbf{X} \in A) &= \sum_{i=1}^m \mathbb{P}(\mathbf{X} \in A | \, \mu = \mu_i) \mathbb{P}(\mu = \mu_i) \\ &= \sum_{i=1}^m w_i \mu_i^{\times n}(A). \end{aligned}$$

The second equality follows from Lemma 3.10 in Kallenberg (2002). So the probability law of $\mathbf{X}$ is

$$\sum_{i=1}^{m} w_i \mu_i^{\times n}. \tag{1}$$

We want to view the probability law of $\mathbf{X}$ as a function of $\mathscr{P}$ in a mathematically rigorous way, which requires a bit of technical buildup. Let $\mathcal{V}$ be a vector space. We will now construct a version of the integral for $\mathcal{V}$-valued functions over $\mathcal{D}$. Let $\mathscr{Q} \in \Delta(\mathcal{D})$. From the definition of $\Delta(\mathcal{D})$ it follows that $\mathscr{Q}$ admits the representation

$$\mathscr{Q} = \sum_{i=1}^{r} \delta_{\mu_i} \alpha_i.$$

From the well-ordering principle there must exist some representation with minimal r and we define r as the *order* of $\mathscr{Q}$. We can show that the representation of any $\mathscr{Q} \in \Delta(\mathcal{D})$ is unique up to permutation of its indices.

Definition 1 *We call $\mathscr{P}$ a* mixture of measures *if it is a probability measure in $\Delta(\mathcal{D})$. We will say that $\mathscr{P}$ has m* mixture components *if it has order m.*

Lemma 2 *Let $\mathscr{Q} \in \Delta(\mathcal{D})$ and admit minimal representations $\mathscr{Q} = \sum_{i=1}^{r} \delta_{\mu_i} \alpha_i = \sum_{i=1}^{r} \delta_{\mu_i'} \alpha_i'$. There exists some permutation $\psi : [r] \to [r]$ such that $\mu_{\psi(i)} = \mu_i'$ and $\alpha_{\psi(i)} = \alpha_i'$ for all i.*

Proof Because both representations are minimal it follows that $\alpha_i' \neq 0$ for all i and $\mu_i' \neq \mu_j'$ for all $i \neq j$. From this we know $\mathscr{Q}(\{\mu_i'\}) \neq 0$ for all i. Because $\mathscr{Q}(\{\mu_i'\}) \neq 0$ for all i it follows that for any i there exists some j such that $\mu_i' = \mu_j$. Let $\psi : [r] \to [r]$ be a function satisfying $\mu_i' = \mu_{\psi(i)}$. Because the elements $\mu_1, \ldots, \mu_r$ are also distinct ψ must be injective and thus a permutation. Again from this distinctness we get that, for all i, $\mathscr{Q}(\{\mu_i'\}) = \alpha_i' = \alpha_{\psi(i)}$ and we are done. ■

Henceforth when we define an element of $\Delta(\mathcal{D})$ with a summation we will assume that the summation is a minimal representation. Any minimal representation of a mixture of measures $\mathscr{P}$ with m components satisfies $\mathscr{P} = \sum_{i=1}^{m} w_i \delta_{\mu_i}$ with $w_i > 0$ for all i and $\sum_{i=1}^{m} w_i = 1$. So any mixture of measures is a convex combination of Dirac measures at elements in $\mathcal{D}$.

For a function $f : \mathcal{D} \to \mathcal{V}$ define

$$\int f(\mu) d\mathscr{Q}(\mu) = \sum_{i=1}^{r} \alpha_i f(\mu_i),$$

where $\sum_{i=1}^{r} \delta_{\mu_i} \alpha_i$ is a minimal representation of $\mathscr{Q}$. This integral is well defined as a consequence of Lemma 2.

For a σ-algebra (Q, Σ) we define $\mathcal{M}(Q, \Sigma)$ as the space of all finite signed measures over that space. Let $\lambda_n : \mathcal{M}(\Omega, \mathcal{F}) \to \mathcal{M}(\Omega^{\times n}, \mathcal{F}^{\times n}); \mu \mapsto \mu^{\times n}$. We introduce the operator $V_n : \Delta(\mathcal{D}) \to \mathcal{M}(\Omega^{\times n}, \mathcal{F}^{\times n})$

$$V_n(\mathscr{Q}) = \int \lambda_n(\mu) d\mathscr{Q}(\mu) = \int \mu^{\times n} d\mathscr{Q}(\mu).$$

For a minimal representation $\mathscr{Q} = \sum_{i=1}^{r} \delta_{\mu_i} \alpha_i$, we have

$$V_n(\mathscr{Q}) = \sum_{i=1}^{r} \mu_i^{\times n} \alpha_i.$$

From this definition we have that $V_n(\mathscr{P})$ is simply the law of $\mathbf{X}$ which we derived earlier. Two mixtures of measures are different if they admit a different measure over $\mathcal{D}$.

Definition 3 *We call a mixture of measures, $\mathscr{P}$,* n-identifiable *if there does not exist a different mixture of measures $\mathscr{P}'$, with order no greater than the order of $\mathscr{P}$, such that $V_n(\mathscr{P}) = V_n(\mathscr{P}')$.*

Definition 3 is the central object of interest in this paper. Given a mixture of measures, $\mathscr{P} = \sum_{i=1}^{m} w_i \delta_{\mu_i}$ then $V_n(\mathscr{P})$ is equal to $\sum_{i=1}^{m} w_i \mu_i^{\times n}$, the measure from which $\mathbf{X}$ is drawn. In topic modelling $\mathbf{X}$ would be the samples from a single document and in transfer learning it would be one of the several collections of training samples. If $\mathscr{P}$ is not n-identifiable then we know that there exists a mixture of measures which is no more complex (in terms of number of mixture components) than $\mathscr{P}$ which is not discernible from $\mathscr{P}$ given the data. Practically speaking this means we need more samples in each random group $\mathbf{X}$ in order for the full richness of $\mathscr{P}$ to be manifested in $\mathbf{X}$.

4. Results

Our primary result gives us a bound on the n-identifiability of all mixtures of measures with m or fewer components. We also show that this bound is tight.

Theorem 4 *Let $(\Omega, \mathcal{F})$ be a measurable space. Mixtures of measures with m components are $(2m-1)$-identifiable.*

Theorem 5 *Let $(\Omega, \mathcal{F})$ be a measurable space with $\mathcal{F} \neq \{\emptyset, \Omega\}$. For all m, there exists a mixture of measures with m components which is not $(2m-2)$-identifiable.*

Unsurprisingly, if a mixture of measures is n-identifiable then it is q-identifiable for all $q > n$. Likewise if a mixture of measures is not n-identifiable then it is not q-identifiable for $q < n$. Thus identifiability is, in some sense, monotonic.

Lemma 6 *If a mixture of measures is n-identifiable then it is q-identifiable for all $q > n$.*

Proof We will proceed by contradiction. Let $\mathscr{P} = \sum_{i=1}^{l} a_i \delta_{\mu_i}$ be n-identifiable, let $\mathscr{P}' = \sum_{j=1}^{r} b_j \delta_{\nu_j}$ be a different mixture of measures with $r \leq l$ and

$$\sum_{i=1}^{l} a_i \mu_i^{\times q} = \sum_{j=1}^{r} b_j \nu_j^{\times q}$$

for some $q > n$. Let $A \in \mathcal{F}^{\times n}$ be arbitrary. We have

$$\sum_{i=1}^{l} a_i \mu_i^{\times q} \quad = \quad \sum_{j=1}^{r} b_j \nu_j^{\times q}$$

$$\begin{aligned} \Rightarrow \sum_{i=1}^{l} a_i \mu_i^{\times q} \left(A \times \Omega^{\times q-n}\right) &= \sum_{j=1}^{r} b_j \nu_j^{\times q} \left(A \times \Omega^{\times q-n}\right) \\ \Rightarrow \sum_{i=1}^{l} a_i \mu_i^{\times n} \left(A\right) &= \sum_{j=1}^{r} b_j \nu_j^{\times n} \left(A\right). \end{aligned}$$

This implies that $\mathscr{P}$ is not n-identifiable, a contradiction. ■

Lemma 7 *If a mixture of measures is not n-identifiable then it is not q-identifiable for any $q < n$.*

Proof Let a mixture of measures $\mathscr{P} = \sum_{i=1}^{l} a_i \delta_{\mu_i}$ not be n-identifiable. It follows that there exists a different mixture of measures $\mathscr{P}' = \sum_{j=1}^{r} b_j \delta_{\nu_j}$, with $r \leq l$, such that

$$\sum_{i=1}^{l} a_i \mu_i^{\times n} = \sum_{j=1}^{r} b_j \nu_j^{\times n}.$$

Let $A \in \mathcal{F}^{\times q}$ be arbitrary, we have

$$\begin{aligned} \sum_{i=1}^{l} a_i \mu_i^{\times n} \left(A \times \Omega^{\times n-q}\right) &= \sum_{j=1}^{r} b_j \nu_j^{\times n} \left(A \times \Omega^{\times n-q}\right) \\ \Rightarrow \sum_{i=1}^{l} a_i \mu_i^{\times q} \left(A\right) &= \sum_{j=1}^{r} b_j \nu_j^{\times q} \left(A\right) \end{aligned}$$

and therefore $\mathscr{P}$ is not q-identifiable. ■

Viewed alternatively these results say that $n = 2m-1$ is the smallest value for which V_n is injective over the set of all minimal mixtures of measures with m or fewer components.

5. Tensor Products of Hilbert Spaces

Our proofs will rely heavily on the geometry of tensor products of Hilbert spaces which we will introduce in this section.

5.1. Overview of Tensor Products

First we introduce tensor products of Hilbert spaces. To our knowledge there does not exist a rigorous construction of the tensor product Hilbert space which is both succinct and intuitive. Because of this we will simply state some basic facts about tensor products of Hilbert spaces and hopefully instill some intuition for the uninitiated by way of example. A through treatment of tensor products of Hilbert spaces can be found in Kadison and Ringrose (1983).

Let H and H' be Hilbert spaces. From these two Hilbert spaces the "simple tensors" are elements of the form $h \otimes h'$ with $h \in H$ and $h' \in H'$. We can treat the simple tensors as being the basis for some inner product space H_0, with the inner product of simple tensors satisfying

$$\left\langle h_1 \otimes h_1', h_2 \otimes h_2' \right\rangle = \left\langle h_1, h_2 \right\rangle \left\langle h_1', h_2' \right\rangle.$$

The tensor product of H and H' is the completion of H_0 and is denoted $H \otimes H'$. To avoid potential confusion we note that notation just described is standard in operator theory literature. In some literature our definition of H_0 is denoted as $H \otimes H'$ and our definition of $H \otimes H'$ is denoted $H \widehat{\otimes} H'$.

As an illustrative example we consider the tensor product $L^2(\mathbb{R}) \otimes L^2(\mathbb{R})$. It can be shown that there exists an isomorphism between $L^2(\mathbb{R}) \otimes L^2(\mathbb{R})$ and $L^2(\mathbb{R}^2)$ which maps the simple tensors to separable functions, $f \otimes f' \mapsto f(\cdot)f'(\cdot)$. We can demonstrate this isomorphism with a simple example. Let $f, g, f', g' \in L^2(\mathbb{R})$. Taking the $L^2(\mathbb{R}^2)$ inner product of $f(\cdot)f'(\cdot)$ and $g(\cdot)g'(\cdot)$ gives us

$$\begin{aligned}\int\int \left(f(x)f'(y)\right)\left(g(x)g'(y)\right)dxdy &= \int f(x)g(x)dx \int f'(y)g'(y)dy \\ &= \langle f, g\rangle \langle f', g'\rangle \\ &= \langle f \otimes f', g \otimes g'\rangle .\end{aligned}$$

Beyond tensor product we will need to define tensor power. To begin we will first show that tensor products are, in some sense, associative. Let H_1, H_2, H_3 be Hilbert spaces. Proposition 2.6.5 in Kadison and Ringrose (1983) states that there is a unique unitary operator, $U : (H_1 \otimes H_2) \otimes H_3 \to H_1 \otimes (H_2 \otimes H_3)$, which satisfies the following for all $h_1 \in H_1, h_2 \in H_2, h_3 \in H_3$,

$$U((h_1 \otimes h_2) \otimes h_3) = h_1 \otimes (h_2 \otimes h_3).$$

This implies that for any collection of Hilbert spaces, $H_1, \ldots, H_n$, the Hilbert space $H_1 \otimes \cdots \otimes H_n$ is defined unambiguously regardless of how we decide to associate the products. In the space $H_1 \otimes \cdots \otimes H_n$ we define a simple tensor as a vector of the form $h_1 \otimes \cdots \otimes h_n$ with $h_i \in H_i$. In Kadison and Ringrose (1983) it is shown that $H_1 \otimes \cdots \otimes H_n$ is the closure of the span of these simple tensors. To conclude this primer on tensor products we introduce the following notation. For a Hilbert space H we denote $H^{\otimes n} = \underbrace{H \otimes H \otimes \cdots \otimes H}_{\text{n times}}$ and for $h \in H$, $h^{\otimes n} = \underbrace{h \otimes h \otimes \cdots \otimes h}_{\text{n times}}$.

5.2. Some Results for Tensor Product Spaces

We will derive state technical results which will be useful for the rest of the paper. These lemmas are similar to or are straightforward extensions of previous results which we needed to modify for our particular purposes. Let $(\Psi, \mathcal{G}, \mu)$ be a σ-finite measure space. We have the following lemma which connects the L^2 space of products of measures to the tensor products of the L^2 space for each measure. The proof of this lemma is straightforward but technical and can be found in the appendix.

Lemma 8 *There exists a unitary transform $U : L^2(\Psi, \mathcal{G}, \mu)^{\otimes n} \to L^2(\Psi^{\times n}, \mathcal{G}^{\times n}, \mu^{\times n})$ such that, for all $f_1, \ldots, f_n \in L^2(\Psi, \mathcal{G}, \mu)$, $U(f_1 \otimes \cdots \otimes f_n) = f_1(\cdot) \cdots f_n(\cdot)$.*

The following lemma used in the proof of Lemma 8 as well as the proof of Theorem 5. The proof of this lemma is also not particularly interesting and can be found in the appendix.

Lemma 9 *Let $H_1, \ldots, H_n, H_1', \ldots, H_n'$ be a collection of Hilbert spaces and $U_1, \ldots, U_n$ a collection of unitary operators with $U_i : H_i \to H_i'$ for all i. There exists a unitary operator $U : H_1 \otimes \cdots \otimes H_n \to H_1' \otimes \cdots \otimes H_n'$ satisfying $U(h_1 \otimes \cdots \otimes h_n) = U_1(h_1) \otimes \cdots \otimes U_n(h_n)$ for all $h_1 \in H_1, \ldots, h_n \in H_n$.*

Lemma 10 *Let $n > 1$ and let $h_1, \ldots, h_n$ be elements of a Hilbert space such that no elements are zero and no pairs of elements are collinear. Then $h_1^{\otimes n-1}, \ldots, h_n^{\otimes n-1}$ are linearly independent.*

A statement of this lemma for $\mathbb{R}^d$ can be found in Comon et al. (2008). We present our own proof for the Hilbert space setting.

Proof We will proceed by induction. For $n = 2$ the lemma clearly holds. Suppose the lemma holds for $n-1$ and let $h_1, \ldots, h_n$ satisfy the assumptions in the lemma statement. Let $\alpha_1, \ldots, \alpha_n$ satisfy

$$\sum_{i=1}^{n} h_i^{\otimes n-1} \alpha_i = 0. \tag{2}$$

To finish the proof we will show that α_1 must be zero which can be generalized to any α_i without loss of generality. Let H_1 and H_2 be Hilbert spaces and let $\mathscr{HS}(H_1, H_2)$ be the space of Hilbert-Schmidt operators from H_1 to H_2. Hilbert-Schmidt operators are a closed subspace of bounded linear operators. Proposition 2.6.9 in Kadison and Ringrose (1983) states that for a pair of Hilbert spaces H_1, H_2 there exists an unitary operator $U : H_1 \otimes H_2 \to \mathscr{HS}(H_1, H_2)$ such that $U(g_1 \otimes g_2) = g_1 \langle g_2, \cdot \rangle$. Applying this operator to (2) we get

$$\sum_{i=1}^{n} h_i^{\otimes n-2} \langle h_i, \cdot \rangle \alpha_i = 0. \tag{3}$$

Because h_1 and h_n are linearly independent we can choose z such that $\langle h_1, z \rangle \neq 0$ and $z \perp h_n$. Plugging z into (3) yields

$$\sum_{i=1}^{n-1} h_i^{\otimes n-2} \langle h_i, z \rangle \alpha_i = 0$$

and therefore $\alpha_1 = 0$ by the inductive hypothesis. ■

6. Proofs of Theorems

With the tools developed in the previous sections we can now prove our theorems. First we introduce one additional piece of notation. For a function p on a domain $\mathcal{X}$ we define $p^{\times k}$ as simply the product of the function k times on the domain $\mathcal{X}^{\times k}$, $\underbrace{p(\cdot) \cdots p(\cdot)}_{\text{k times}}$. For a measure the notation continues to denote the standard product measure.

Finally will need the following technical lemma to connect the product of Radon-Nikodym derivatives to product measures. The proof is straightforward and can be found in the appendix.

Lemma 11 *Let $(\Psi, \mathcal{G})$ be a measurable space, η and γ a pair of bounded measures on that space, and f a nonnegative function in $L^1(\gamma)$ such that, for all $A \in \mathcal{G}$, $\eta(A) = \int_A f d\gamma$. Then for all n, for all $B \in \mathcal{G}^{\times n}$ we have*

$$\eta^{\times n}(B) = \int_B f^{\times n} d\gamma^{\times n}.$$

Proof of Theorem 4 We will proceed by contradiction. Suppose there exist two different mixtures of measures $\mathscr{P} = \sum_{i=1}^{l} \delta_{\mu_i} a_i \neq \mathscr{P}' = \sum_{j=1}^{m} \delta_{\nu_j} b_j$, such that

$$\sum_{i=1}^{l} a_i {\mu_i}^{\times 2m-1} = \sum_{j=1}^{m} b_j \nu_j^{\times 2m-1}$$

and $l \leq m$. From our assumption on representation we know $\mu_i \neq \mu_j$ for all $i \neq j$ and similarly for $\nu_1, \ldots, \nu_m$. We will also assume that $\mu_i \neq \nu_j$ for all i, j. Were this not true we could simply subtract the smaller of the common terms from both sides of (4) and normalize to yield another pair of distinct mixtures of measures with fewer components and no shared terms, $\mathscr{Q}$ and $\mathscr{Q}'$. Let $\mathscr{Q}$ have m' components and $\mathscr{Q}'$ have l' with $m' \geq l'$. If $m \neq m'$ then we can apply Lemma 7 to give us $V_{2m'-1}(\mathscr{Q}) = V_{2m'-1}(\mathscr{Q}')$ and proceed as usual.

Let $\xi = \sum_{i=1}^{l} \mu_i + \sum_{j=1}^{m} \nu_j$. Clearly ξ dominates μ_i and ν_j for all i, j so we can define Radon-Nikodym derivatives $p_i = \frac{d\mu_i}{d\xi}$, $q_j = \frac{d\nu_j}{d\xi}$ which are in $L^1(\Omega, \mathcal{F}, \xi)$. We can assert that these derivatives are everywhere nonnegative without issue. Clearly no two of these derivatives are equal. If one of the derivatives were a scalar multiple of another, for example $p_1 = \alpha p_2$ for some $\alpha \neq 1$, it would imply

$$\mu_1(\Omega) = \int_{\Omega} p_1 d\xi = \int \alpha p_2 d\xi = \alpha.$$

This is not true so no pair of these derivatives are collinear.

Lemma 11 tells us that, for any $R \in \mathcal{F}^{\times 2m-1}$ we have

$$\begin{aligned} \int_R \sum_{i=1}^{l} a_i p_i^{\times 2m-1} d\xi^{\times 2m-1} &= \sum_{i=1}^{l} a_i \mu_i^{\times 2m-1}(R) \\ &= \sum_{j=1}^{m} b_j \nu_j^{\times 2m-1}(R) \\ &= \int_R \sum_{j=1}^{m} b_j q_j^{\times 2m-1} d\xi^{\times 2m-1}. \end{aligned}$$

Therefore

$$\sum_{i=1}^{l} a_i p_i^{\times 2m-1} = \sum_{j=1}^{m} b_j q_j^{\times 2m-1} \tag{4}$$

$\xi^{\times 2m-1}$-almost everywhere (Proposition 2.23 in Folland (1999)). We will now show for all i, j that $p_i \in L^2(\Omega, \mathcal{F}, \xi)$ and $q_j \in L^2(\Omega, \mathcal{F}, \xi)$. We will argue this for p_1 which will clearly generalize to the other elements. First we will show that $p_1 \leq 1$ ξ-almost everywhere. Suppose this were not true and that there exists $A \in \mathcal{F}$ with $\xi(A) > 0$ and $p_1(A) > 1$. Now we would have

$$\mu_1(A) = \int_A p_1 d\xi > \int_A 1 d\xi = \xi(A) = \sum_{i=1}^{l} \mu_i(A) + \sum_{j=1}^{m} \nu_j(A) \geq \mu_1(A)$$

a contradiction. Evaluating directly we get

$$\begin{aligned}\int p_1(\omega)^2 d\xi\left(\omega\right) &\leq \int 1 d\xi\left(\omega\right)\\ &= \xi\left(\Omega\right)\\ &= l+m,\end{aligned}$$

so $p_1 \in L^2\left(\Omega, \mathcal{F}, \xi\right)$. Applying the U^{-1} operator from Lemma 8 to (4) yields

$$\sum_{i=1}^{l} a_i p_1^{\otimes 2m-1} = \sum_{j=1}^{m} b_j q_j^{\otimes 2m-1}.$$

Since $l+m \leq 2m$ Lemma 10 states that $p_1^{\otimes 2m-1}, \ldots, p_l^{\otimes 2m-1}, q_1^{\otimes 2m-1}, \ldots, q_m^{\otimes 2m-1}$ are all linearly independent and thus $a_i = 0$ and $b_j = 0$ for all i, j, a contradiction. ■

Proof of Theorem 5 To prove this theorem we will construct a pair of different mixture of measures, $\mathscr{P} \neq \mathscr{P}'$ which both contain m components and satisfy $V_{2m-2}\left(\mathscr{P}\right) = V_{2m-2}\left(\mathscr{P}'\right)$.

From our definition of $(\Omega, \mathcal{F})$ we know there exists $F \in \mathcal{F}$ such that F, F^C are nonempty. Let $f \in F$ and $f' \in F^C$. It follows that $\delta_f \neq \delta_{f'}$ are different probability measures on $(\Omega, \mathcal{F})$. Because δ_f and $\delta_{f'}$ are dominated by $\xi = \delta_f + \delta_{f'}$ we know that there exists a pair of measurable functions p, p' such that, for all A, $\delta_f\left(A\right) = \int_A p d\xi$ and $\delta_{f'}\left(A\right) = \int_A p' d\xi$. We can assert that p and p' are nonnegative without issue.

From the same argument we used in the proof of Theorem 4 we know $p, p' \in L^2\left(\Omega, \mathcal{F}, \xi\right)$. Let H_2 be the Hilbert space generated from the span of p, p'. Let $(\varepsilon_i)_{i=1}^{2m}$ be $2m$ distinct elements of $[0,1]$ and let $(p_i)_{i=1}^{2m}$ be elements of $L^1(\Omega, \mathcal{F}, \xi)$ with $p_i = \varepsilon_i p + (1-\varepsilon_i) p'$. Clearly p_i is a pdf over ξ for all i and there are no pairs in this collection which are collinear. Let H_2 be the Hilbert space generated from the span of p and p'. Since H_2 is isomorphic to $\mathbb{R}^2$ there exists a unitary operator $U : H_2 \to \mathbb{R}^2$. From Lemma 9 there exists a unitary operator $U_{2m-2} : H_2^{\otimes 2m-2} \to \mathbb{R}^{2^{\otimes 2m-2}}$ with $U_{2m-2}\left(h_1 \otimes \cdots \otimes h_{2m-2}\right) = U(h_1) \otimes \cdots \otimes U(h_{2m-2})$. Because U is unitary the set $U_{2m-2}\left(\text{span}\left(\left\{h^{\otimes 2m-2} : h \in H_2\right\}\right)\right)$ maps exactly to the set $\text{span}\left(x^{\otimes 2m-2} : x \in \mathbb{R}^2\right)$. An order r tensor, $A_{i_1,\ldots,i_r}$, is *symmetric* if $A_{\psi(i_1),\ldots,\psi(i_r)} = A_{i_1,\ldots,i_r}$ for any $i_1, \ldots, i_r$ and permutation ψ. A consequence of Lemma 4.2 in Comon et al. (2008) is that $\text{span}\left(\left\{x^{\otimes 2m-2} : x \in \mathbb{R}^2\right\}\right) \subset S^{2m-2}(\mathbb{C}^2)$ is exactly the space of all symmetric order $2m-2$ tensors over $\mathbb{C}^2$.

From Proposition 3.4 in Comon et al. (2008) it follows that the dimension of $S^{2m-2}\left(\mathbb{C}^2\right)$ is $\begin{pmatrix} 2+2m-2-1 \\ 2m-2 \end{pmatrix} = 2m-1$. From this we get that $\dim\left(\text{span}\left(\left\{h^{\otimes 2m-2} : h \in H_2\right\}\right)\right) \leq 2m-1$.

The bound on the dimension of $\text{span}\left(\left\{h^{\otimes 2m-2} : h \in H_2\right\}\right)$ implies that $\left(p_i^{\otimes 2m-2}\right)_{i=1}^{2m}$ are linearly dependent. Conversely Lemma 10 implies that removing a single vector from $\left(p_i^{\otimes 2m-2}\right)_{i=1}^{2m}$ yields a set of vectors which are linearly independent. It follows that there exists $(\alpha_i)_{i=1}^{2m}$ with $\alpha_i \neq 0$ for all i and

$$\sum_{i=1}^{2m} \alpha_i p_i^{\otimes 2m-2} = 0.$$

Without loss of generality we will assume that $\alpha_i < 0$ for $i \in [k]$ with $k \leq m$. From this we have

$$\sum_{i=1}^{k} -\alpha_i p_i^{\otimes 2m-2} = \sum_{j=k+1}^{2m} \alpha_j p_j^{\otimes 2m-2}. \tag{5}$$

From Lemma 8 we have

$$\sum_{i=1}^{k} -\alpha_i p_i^{\times 2m-2} = \sum_{j=k+1}^{2m} \alpha_j p_j^{\times 2m-2}$$

and thus

$$\begin{aligned} \int \sum_{i=1}^{k} -\alpha_i p_i^{\times 2m-2} d\xi^{\times 2m-2} &= \int \sum_{j=k+1}^{2m} \alpha_j p_j^{\times 2m-2} d\xi^{\times 2m-2} \\ \Rightarrow \sum_{i=1}^{k} -\alpha_i &= \sum_{j=k+1}^{2m} \alpha_j. \end{aligned}$$

Let $r = \sum_{i=1}^{k} -\alpha_i$. We know $r > 0$ so dividing both sides of (5) by r gives us

$$\sum_{i=1}^{k} -\frac{\alpha_i}{r} p_i^{\otimes 2m-2} = \sum_{j=k+1}^{2m} \frac{\alpha_j}{r} p_j^{\otimes 2m-2}$$

and the left and the right side are convex combinations. Let $(\beta_i)_{i=1}^{2m}$ positive numbers with $\beta_i = \frac{-\alpha_i}{r}$ for $i \in \{1, \ldots, k\}$ and $\beta_j = \frac{\alpha_j}{r}$ for $j \in \{k+1, \ldots, 2m\}$. This gives us

$$\sum_{i=1}^{k} \beta_i p_i^{\otimes 2m-2} = \sum_{j=k+1}^{2m} \beta_j p_j^{\otimes 2m-2}.$$

It follows that

$$\sum_{i=1}^{k} \beta_i p_i^{\otimes m-1} \otimes p_i^{\otimes m-1} = \sum_{j=k+1}^{2m} \beta_j p_j^{\otimes m-1} \otimes p_i^{\otimes m-1}.$$

We will now show that $k = m$. Suppose $k < m$. Then $p_1^{\otimes m-1}, \ldots, p_{k+1}^{\otimes m-1}$ are linearly independent. From this we know that there exists z such that $z \perp p_i^{\otimes m-1}$ for $i \in [k]$ but z is not orthogonal to $p_{k+1}^{\otimes m-1}$. Using this vector we have

$$\begin{aligned} \left\langle \sum_{i=1}^{k} \beta_i p_i^{\otimes 2m-1}, z \otimes z \right\rangle &= \sum_{i=1}^{k} \beta_i \left\langle z, p_i^{\otimes m-1} \right\rangle \left\langle z, p_i^{\otimes m-1} \right\rangle \\ &= 0 \end{aligned}$$

but

$$\left\langle \sum_{i=k+1}^{2m} \beta_i p_i^{\otimes m-1} \otimes p_i^{\otimes m-1}, z \otimes z \right\rangle = \sum_{i=k+1}^{2m} \beta_i \left\langle p_i^{\otimes m-1}, z \right\rangle \left\langle p_i^{\otimes m-1}, z \right\rangle > 0$$

and thus $k = m$.

Now we have

$$\sum_{i=1}^{m} \beta_i p_i^{\otimes 2m-2} = \sum_{j=m+1}^{2m} \beta_j p_j^{\otimes 2m-2}.$$

Applying Lemma 8 we get that

$$\sum_{i=1}^{m} \beta_i p_i^{\times 2m-2} = \sum_{j=m+1}^{2m} \beta_j p_j^{\times 2m-2}.$$

From Lemma 11 we have,

$$\sum_{i=1}^{m} \beta_i \left(\varepsilon_i \delta_f + (1-\varepsilon_i)\, \delta_{f'}\right)^{\times 2m-2} = \sum_{j=m+1}^{2m} \beta_j \left(\varepsilon_j \delta_f + (1-\varepsilon_j)\, \delta_{f'}\right)^{\times 2m-2}.$$

Setting $\mu_i = \left(\varepsilon_i \delta_f + (1-\varepsilon_i)\, \delta_{f'}\right)$ yields

$$\sum_{i=1}^{m} \beta_i \mu_i^{\times 2m-2} = \sum_{j=m+1}^{2m} \beta_j \mu_j^{\times 2m-2}.$$

Thus setting $\mathscr{P} = \sum_{i=1}^{m} \beta_i \delta_{\mu_i}$ and $\mathscr{P}' = \sum_{j=m+1}^{2m} \beta_j \delta_{\mu_j}$ gives us $V_{2m-2}(\mathscr{P}) = V_{2m-2}(\mathscr{P}')$ and $\mathscr{P} \neq \mathscr{P}'$ by construction. ■

6.1. Discussion of the Proof of Theorem 5

In the previous proof we could have replaced $\delta_f, \delta_{f'}$ with any distinct pair of probability measures on $(\Omega, \mathcal{F})$. Thus the pair $\mathscr{P}, \mathscr{P}'$ are not pathological because of some property of each individual mixture component, but because of geometry of the mixture components considered as a whole. The measures $\mu_1, \ldots, \mu_{2n}$ are a convex combinations of δ_f and $\delta_{f'}$ and therefore lie in a one dimensional affine subspace of $\Delta(\mathcal{D})$. The space of Bernoulli measures similarly lie in a subspace between two measures, the point mass at 0 and the point mass at 1. Given a mixture of Bernoulli distributions, the sum of iid samples of Bernoulli random variables is a binomial distribution. We can draw a connection between our result and the identifiability of mixtures of binomial distributions.

Consider $\mathscr{P}$ as mixture of m Bernoulli distributions with parameters $\lambda_1, \ldots, \lambda_m$ and weights $w_1, \ldots w_m$. Suppose we have n samples in each random group. If we let Y_i be the sum of the random group $\mathbf{X}_i$ then the probability law of Y_i is a mixture of binomial random variables. Let $p(\lambda, n)$ be the distribution of a Bernoulli random variable with parameters n and λ. Specifically we have that the distribution of $Y_i = \sum_{i=1}^{m} w_i p(\lambda_i, n)$. In Blischke (1964) it was shown that $n \geq 2m-1$ is a necessary and sufficient condition for the identifiability of the parameters $\lambda_1, \ldots, \lambda_m$ from the samples Y_i. We find these similarities provoking but are not prepared to make more precise connections at this time.

7. Conclusion

In this paper we have proven a fundamental bound on the identifiability of mixture models in a nonparametric setting. Any mixture with m components is identifiable with groups of samples containing $2m - 1$ samples from the same latent probability measure. We show that this bound is tight by constructing a mixture of m probability measures which is not identifiable with groups of samples containing $2m - 2$. These results hold for any mixture over any domain with at least two elements.

References

Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *Ann. Statist.*, 37(6A):3099–3132, 12 2009. doi: 10.1214/09-AOS689. URL http://dx.doi.org/10.1214/09-AOS689.

Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15: 2773–2832, 2014. URL http://jmlr.org/papers/v15/anandkumar14b.html.

Joseph Anderson, Mikhail Belkin, Navin Goyal, Luis Rademacher, and James Voss. The more, the merrier: the blessing of dimensionality for learning large gaussian mixtures. In *Proceedings of The 27th Conference on Learning Theory*, pages 1135–1164, 2014.

Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization – provably. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 145–162, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1245-5. doi: 10.1145/2213977.2213994. URL http://doi.acm.org/10.1145/2213977.2213994.

Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. In John Shawe-Taylor, Richard S. Zemel, Peter L. Bartlett, Fernando C. N. Pereira, and Kilian Q. Weinberger, editors, *NIPS*, pages 2178–2186, 2011. URL http://dblp.uni-trier.de/db/conf/nips/nips2011.html#BlanchardLS11.

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, March 2003. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=944919.944937.

W. R. Blischke. Estimating the parameters of mixtures of binomial distributions. *Journal of the American Statistical Association*, 59(306):pp. 510–528, 1964. ISSN 01621459. URL http://www.jstor.org/stable/2283005.

Pierre Comon, Gene Golub, Lek-Heng Lim, and Bernard Mourrain. Symmetric tensors and symmetric tensor rank. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1254–1279, 2008. doi: 10.1137/060661569. URL http://dx.doi.org/10.1137/060661569.

Sanjoy Dasgupta and Leonard Schulman. A probabilistic analysis of em for mixtures of separated, spherical gaussians. *J. Mach. Learn. Res.*, 8:203–226, May 2007. ISSN 1532-4435. URL http://portal.acm.org/citation.cfm?id=1248659.1248666.

Gerald B. Folland. *Real analysis: modern techniques and their applications*. Pure and applied mathematics. Wiley, 1999. ISBN 9780471317166. URL http://books.google.com/books?id=uPkYAQAAIAAJ.

T. P. Hettmansperger and Hoben Thomas. Almost nonparametric inference for repeated measures in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(4):811–825, 2000. ISSN 1467-9868. doi: 10.1111/1467-9868.00266. URL http://dx.doi.org/10.1111/1467-9868.00266.

R.V. Kadison and J.R. Ringrose. *Fundamentals of the theory of operator algebras. V1: Elementary theory*. Pure and Applied Mathematics. Elsevier Science, 1983. ISBN 9780080874166. URL https://books.google.com/books?id=JbxgKOwu2McC.

Olav Kallenberg. *Foundations of modern probability*. Probability and its applications. Springer, New York, Berlin,, Paris, 2002. ISBN 0-387-95313-2. URL http://opac.inria.fr/record=b1098179. Sur la 4e de couv. : This new edition contains four new chapters as well as numerous improvements throughout the text.

Joseph B. Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra and its Applications*, 18(2):95 – 138, 1977. ISSN 0024-3795.

Andreas Maurer, Massi Pontil, and Bernardino Romera-paredes. Sparse coding for multitask and transfer learning. In Sanjoy Dasgupta and David Mcallester, editors, *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, volume 28, pages 343–351. JMLR Workshop and Conference Proceedings, May 2013. URL http://jmlr.csail.mit.edu/proceedings/papers/v28/maurer13.pdf.

Krikamol Muandet and Bernhard Schölkopf. One-class support measure machines for group anomaly detection. *CoRR*, abs/1303.0309, 2013. URL http://dblp.uni-trier.de/db/journals/corr/corr1303.html#abs-1303-0309.

Barnabás Póczos, Aarti Singh, Alessandro Rinaldo, and Larry A. Wasserman. Distribution-free distribution regression. In *AISTATS*, volume 31 of *JMLR Proceedings*, pages 507–515. JMLR.org, 2013. URL http://dblp.uni-trier.de/db/conf/aistats/aistats2013.html#PoczosSRW13.

Yuval Rabani, Leonard J. Schulman, and Chaitanya Swamy. Learning mixtures of arbitrary distributions over large discrete domains. *ArXiv e-prints*, 2013. URL http://arxiv.org/abs/1212.1527.

Z. Szabo, B. Sriperumbudur, B. Poczos, and A. Gretton. Learning Theory for Distribution Regression. *ArXiv e-prints*, November 2014.

Appendix A. Additional Proofs

Proof of Lemma 8 Example 2.6.11 in Kadison and Ringrose (1983) states that for any two σ-finite measure spaces $(S, \mathscr{S}, m), (S', \mathscr{S}', m')$ there exists a unitary operator $U : L^2(S, \mathscr{S}, m) \otimes L^2(S', \mathscr{S}', m') \to L^2(S \times S', \mathscr{S} \times \mathscr{S}', m \times m')$ such that, for all f, g,

$$U(f \otimes g) = f(\cdot)g(\cdot).$$

Because $(\Psi, \mathcal{G}, \eta)$ is a σ-finite measure space it follows that $(\Psi^{\times m}, \mathcal{G}^{\times m}, \eta^{\times m})$ is a σ-finite measure space for all $m \in \mathbb{N}$. We will now proceed by induction. Clearly the lemma holds for $n = 1$. Suppose the lemma holds for $n - 1$. From the induction hypothesis we know that there exists a unitary transform $U_{n-1} : L^2(\Psi, \mathcal{G}, \eta)^{\otimes n-1} \to L^2\left(\Psi^{\times n-1}, \mathcal{G}^{\times n-1}, \eta^{n-1}\right)$ such that for all simple tensors $f_1 \otimes \cdots \otimes f_{n-1} \mapsto f_1(\cdot) \cdots f_{n-1}(\cdot)$. Combining U_{n-1} with the identity map via Lemma 9 we can construct a unitary operator $T_n : L^2(\Psi, \mathcal{G}, \eta)^{\otimes n-1} \otimes L^2(\Psi, \mathcal{G}, \eta) \to L^2\left(\Psi^{\times n-1}, \mathcal{G}^{\times n-1}, \eta^{n-1}\right) \otimes L^2(\Psi, \mathcal{G}, \eta)$, which maps $f_1 \otimes \cdots \otimes f_{n-1} \otimes f_n \mapsto f_1(\cdot) \cdots f_{n-1}(\cdot) \otimes f_n$

From the aforementioned example there exists a unitary transform $K_n : L^2\left(\Psi^{n-1}, \mathcal{G}^{\times n-1}, \eta^{n-1}\right) \otimes L^2(\Psi, \mathcal{G}, \eta) \to L^2\left(\Psi^{\times n-1} \times \Psi, \mathcal{G}^{\times n-1} \times \mathcal{G}, \eta^{n-1} \times \eta\right)$ which maps $f \otimes f' \mapsto f(\cdot) f'(\cdot)$. Defining $U_n(\cdot) = K_n(T_n(\cdot))$ yields our desired unitary transform. ■

Proof of Lemma 9 Proposition 2.6.12 in Kadison and Ringrose (1983) states that there exists a continuous linear operator $\tilde{U} : H_1 \otimes \cdots \otimes H_n \to H_1' \otimes \cdots \otimes H_n'$ such that $\tilde{U}(h_1 \otimes \cdots \otimes h_n) = U_1(h_1) \otimes \cdots \otimes U_n(h_n)$ for all $h_1 \in H_1, \cdots, h_n \in H_n$. Let $\widehat{H}$ be the set of simple tensors in $H_1 \otimes \cdots \otimes H_n$ and $\widehat{H}'$ be the set of simple tensors in $H_1' \otimes \cdots \otimes H_n'$. Because U_i is surjective for all i, clearly $\tilde{U}(\widehat{H}) = \widehat{H}'$. The linearity of $\tilde{U}$ implies that $\tilde{U}(\operatorname{span}(\widehat{H})) = \operatorname{span}(\widehat{H}')$. Because $\operatorname{span}(\widehat{H}')$ is dense in $H_1' \otimes \cdots \otimes H_n'$ the continuity of $\tilde{U}$ implies that $\tilde{U}(H_1 \otimes \cdots \otimes H_n) = H_1' \otimes \cdots \otimes H_n'$ so $\tilde{U}$ is surjective. All that remains to be shown is that $\tilde{U}$ preserves the inner product. By the continuity of inner product we need only show that $\langle h, g \rangle = \left\langle \tilde{U}(h), \tilde{U}(g) \right\rangle$ for $h, g \in \operatorname{span}(\widehat{H})$. With this in mind let $h_1, \ldots, h_N, g_1, \ldots, g_M \in \widehat{H}$. We have the following

$$\begin{aligned}
\left\langle \tilde{U}\left(\sum_{i=1}^{N} h_i\right), \tilde{U}\left(\sum_{j=1}^{M} g_j\right) \right\rangle &= \left\langle \sum_{i=1}^{N} \tilde{U}(h_i), \sum_{j=1}^{M} \tilde{U}(g_j) \right\rangle \\
&= \sum_{i=1}^{N} \sum_{j=1}^{M} \left\langle \tilde{U}(h_i), \tilde{U}(g_j) \right\rangle \\
&= \sum_{i=1}^{N} \sum_{j=1}^{M} \langle h_i, g_j \rangle \\
&= \left\langle \sum_{i=1}^{N} h_i, \sum_{j=1}^{M} g_j \right\rangle.
\end{aligned}$$

We have now shown that $\tilde{U}$ is unitary which completes our proof. ■

Proof of Lemma 11 The fact that f is positive and integrable implies that the map $S \mapsto \int_S f^{\times n} d\gamma^{\times n}$ is a bounded measure on $(\Psi^{\times n}, \mathcal{G}^{\times n})$ (see Folland (1999) Exercise 2.12).

Let $R = R_1 \times \ldots \times R_n$ be a rectangle in $\mathcal{G}^{\times n}$. Let $\mathbb{1}_S$ be the indicator function for a set S. Integrating over R and using Tonelli's theorem we get

$$\begin{aligned}
\int_R f^{\times n} d\gamma^{\times n} &= \int \mathbb{1}_R f^{\times n} d\gamma^{\times n} \\
&= \int \mathbb{1}_R f^{\times n} d\gamma^{\times n} \\
&= \int \left(\prod_{i=1}^n \mathbb{1}_{R_i}(x_i) \right) \left(\prod_{j=1}^n f(x_j) \right) d\gamma^{\times n}(x_1, \ldots, x_n) \\
&= \int \cdots \int \left(\prod_{i=1}^n \mathbb{1}_{R_i}(x_i) \right) \left(\prod_{j=1}^n f(x_j) \right) d\gamma(x_1) \cdots d\gamma(x_n) \\
&= \int \cdots \int \left(\prod_{i=1}^n \mathbb{1}_{R_i}(x_i) f(x_i) \right) d\gamma(x_1) \cdots d\gamma(x_n) \\
&= \prod_{i=1}^n \left(\int \mathbb{1}_{R_i}(x_i) f(x_i) d\gamma(x_i) \right) \\
&= \prod_{i=1}^n \eta(R_i) \\
&= \eta^{\times n}(R).
\end{aligned}$$

Any product probability measure is uniquely determined by its measure over the rectangles (this is a consequence of Lemma 1.17 in Kallenberg (2002) and the definition of product σ-algebra) therefore, for all $B \in \mathcal{G}^n$,

$$\eta^{\times n}(B) = \int_B f^{\times n} d\gamma^{\times n}.$$

■